NETWORK ADDRESS TRANSLATION IN WINDOWS 2000

> **After reading this chapter and completing the exercises, you will be able to:**
>
> ♦ Explain the differences between Internet Connection Sharing (ICS) and Network Address Translation (NAT)
> ♦ Describe the address translation process
> ♦ Install and configure ICS on Windows 2000 Server or Professional
> ♦ Install and configure NAT on Windows 2000 Server
> ♦ Monitor and manage NAT

Network Address Translation (NAT) is a protocol that provides a way for multiple computers on a network to share a single connection to the Internet via an Internet Service Provider. In Windows 2000, two different services provide access to this protocol and you choose a particular service based on your networking needs. In typical Microsoft style, the names of these services often generate a bit of confusion.

Internet Connection Sharing (ICS), a service that is easy to configure and manage, offers most of the features of the Network Address Translation protocol. However, you cannot control many ICS features. It's more of a "turn it on and watch it run" service. ICS is available in a number of Microsoft operating systems, including Windows 98 Second Edition, Windows Millennium Edition, Windows 2000 Professional, and Windows 2000 Server (or Advanced Server).

NAT runs only on the Windows 2000 Server family and is implemented as a routing protocol within the Routing and Remote Access Service that you learned about in earlier chapters. While it provides many of the same services as ICS, NAT is much more configurable and offers some added features discussed later in this chapter.

> **Note** Confusion often arises because one of the implementations of the Network Address Translation protocol is named NAT. This chapter uses the full name "Network Address Translation," to refer to the protocol itself. The abbreviation "NAT" refers to the implementation of the protocol within RRAS. Keep in mind, however, that other literature and the certification exam may refer to the protocol and the service in either way.

This chapter begins with an overview of address translation and the differences between ICS and NAT. From there, the chapter moves on to the actual configuration and management of both these services.

OVERVIEW

Until recently, most operating systems did not include a way for more than one computer to use a single connection to the Internet. In a typical setup, a computer was configured to use a dial–up connection (like a modem or ISDN adapter) or a persistent connection (like a DSL line or cable modem) to connect to an Internet Service Provider. If you had more than one computer, say on a small home or office network, you were forced to configure a sep–arate connection for each system or purchase a third–party proxy program to allow those computers to share access.

> **Note** **SOHO** is an acronym for Small Office/Home Office. Microsoft regards SOHO networks as the main beneficiaries of ICS and NAT. Though SOHO networks configuration varies a great deal, Microsoft normally considers a SOHO net–work to have one network segment, use peer-to-peer networking, and support TCP/IP. For larger networks, Microsoft generally recommends a separate prod–uct, such as Microsoft Proxy Server, to provide address translation services. In the real world, these definitions really don't mean too much. NAT is often used on large networks quite effectively. However, for the certification exam, you should be aware of the distinctions that Microsoft draws.

With the advent of Windows 98 Second Edition, Microsoft began incorporating a simpli–fied version of the Network Address Translation protocol into the operating system so that no third–party software was required to share Internet connections. They named the service Internet Connection Sharing. Windows Millennium Edition, Windows 2000 Professional, and Windows 2000 Server also come with ICS. In addition, Windows 2000 Server supports the full version of NAT, which offers a good deal more flexibility than ICS.

This overview discusses the Network Address Translation protocol and address sharing in a conceptual fashion. The end of the overview presents the actual differences between these two implementations of the Network Address Translation protocol.

Benefits of Address Sharing

So, why share addresses in the first place? Address sharing really provides three benefits:

- Using address translation instead of routing provides an inherent security benefit. Hosts on the Internet only see the public IP address of the external interface on the computer that provides address translation—not the private IP addresses on the internal network.

- Cost is another big reason to share addresses. It's obviously cheaper to configure one computer with a high-speed Internet connection than to provide one for every computer on your network.

- Simplicity is the third reason to share addresses. Setting up one Internet connection (especially with some of the more complicated connection options out there today) and then sharing that connection is easier than configuring a connection for every computer.

Public and Private Addressing

In Chapter 2, you learned all the gory details of IP addressing, including the different classes of addresses available on the Internet. You also learned that, although you can subnet and supernet your networks in many ways to maximize the efficiency of IP address assignments, only a finite number of IP addresses are available. In addition, the amazing growth of the Internet has greatly strained the capacity of current IP addressing.

In an early attempt to work around this problem, the Internet Network Information Center (InterNIC) and the Internet Assigned Numbers Authority (IANA) designated three network IDs as private networks:

- 10.0.0.0 with a subnet mask of 255.0.0.0. This provides a range of **private addresses** from 10.0.0.1 through 10.255.255.254.

- 172.16.0.0 with a subnet mask of 255.240.0.0. This provides a range of private addresses from 172.16.0.1 through 172.31.255.254.

- 192.168.0.0 with a subnet mask of 255.255.0.0. This provides a range of private addresses from 192.168.0.1 through 192.168.255.254.

No host with any of the addresses in these ranges is ever allowed to transfer information directly to a host on the Internet that has a **public address**. The original intent behind assigning these private address ranges was that they would be used on networks that would not connect to the Internet. You could address your local network, and even subnet it, any way you liked as long as your addresses stayed within the private ranges and did not try to connect to any public hosts.

With NAT, private networks now have a way of transferring information to the Internet, even though they use private addresses. Your ISP only need assign you one public IP address (though NAT can handle multiple public addresses), and NAT translates between the private IP addresses on your network and that public IP address. To the Internet, it looks like

you have one host (the NAT server), even though your private network may have dozens of computers hiding behind that host.

How NAT Works

A NAT server is basically an IP router that translates the IP addresses and TCP/UDP port numbers of packets as those packets are forwarded between the public and private interfaces of the NAT server. This section examines the actual NAT process in more detail.

Static and Dynamic Address Mapping

When NAT receives a packet from a private IP address and translates that packet to look as though it comes from the NAT server's public IP address, this process is called "mapping." Two forms of mapping are available in NAT:

- **Dynamic mappings** are created when users on the private network initiate traffic with a public Internet location. The NAT service automatically translates the IP address and source ports, and adds these mappings to its mapping table. The NAT server refreshes these mappings each time they are used. Dynamic mappings that are not refreshed are removed from the NAT mapping table after a certain amount of time. For TCP connections, the default time is 24 hours. For UDP connections, the default time is one minute.

- **Static mappings** define in advance the mapping of certain addresses and ports instead of letting mapping happen automatically. Although you can create static mappings for outbound traffic, the most common reason to use static mapping is if you want to host some form of Internet service (that is, Web server, FTP server, and so forth.) on a private computer. For hosts on the Internet to reach that server, a static mapping must be defined so that the NAT server knows where to route the incoming requests. You cannot host any Internet services on your private network using dynamic mapping.

NAT Editors

For NAT to translate packets directly between a private and public network, two things must be true:

- The packets must have an IP address in the IP header.
- The packets must have either a TCP or UDP port number in the IP header.

While this works fine for the majority of protocols and applications that send IP traffic (since many of them use TCP or UDP), some do not fulfill these requirements. For example, neither FTP nor PPTP uses TCP or UDP, so NAT could not translate them without a little help.

This help comes in the form of a **NAT editor**, an installable component that modifies packets so that NAT can translate them. Windows 2000 includes built-in NAT editors for the following protocols:

- FTP
- Internet Message Control Protocol (ICMP)
- Point-to-Point Tunneling Protocol (PPTP)
- NetBIOS over TCP/IP (NetBT)

In addition to the built-in NAT editors, the NAT protocol in Windows 2000 includes proxy software for the following protocols:

- H.323, a protocol voice and data transmission
- Direct Play, a protocol used in multiplayer gaming
- LDAP-based Internet Locator Service (ILS) registration, a protocol used by NetMeeting
- Remote Procedure Call (RPC)

It is important to note that the NAT protocol does not at this point support either the Kerberos authentication method used in Windows or the IPSec protocol. Chapter 8 discusses both of these protocols.

DHCP Allocator

Both forms of NAT offered by Windows 2000 (ICS and NAT) can automatically assign IP addresses to computers on the private network using a **DHCP Allocator**, a simplified version of a DHCP server. This works well on small networks, as most clients are set up to receive IP addresses automatically by default.

> You can learn more about using DHCP in Chapter 3 and more about configuring it to work with NAT later in this chapter.

When a client starts, it broadcasts a message looking for DHCP allocation; the NAT server assigns it an IP address and subnet mask on the same subnet using a private addressing range. In addition, the NAT server configures the default gateway and DNS server for clients to be the IP address of the NAT server. Note that there is no WINS server allocation.

As you learn later in the chapter, the DHCP Allocator in ICS is enabled by default and cannot be disabled. Although you can assign static addresses to the other computers on the network if you want, the ICS server always responds to DHCP requests. When using NAT on a Windows 2000 Server, you can disable the DHCP Allocator and either assign static addresses from the NAT server or let another DHCP Server on the network handle requests.

Host Name Resolution

When using the DHCP Allocator, clients are configured to use the NAT server as their primary DNS server. This allows both local and remote host names to be resolved. **DNS proxying** is used to resolve remote host names on the Internet. In this process, a client submits a name resolution request to the NAT server. The NAT server then queries the DNS server specified in its own configuration for the resolution. When it receives a response, it forwards that response to the originating client.

Differences Between NAT and ICS

Since both ICS and NAT use the same protocol to translate addresses, this overview features a combined discussion of their similar features. Table 9-1 shows how each service implements the NAT protocol differently.

Table 9-1   Differences between ICS and NAT

ICS	NAT
Available on Windows 98 Second Edition, Windows Millennium Edition, Windows 2000 Professional, and the Windows 2000 Server family	Available only on the Windows 2000 Server family
Configured in Windows 2000 by checking a single option on the Sharing page of a network adapter	Requires you to use the Routing and Remote Access snap-in for installation and management; provides a lot more configuration options
Allows only one public IP address	Can expose any number of public addresses
Links only one private network to a public network	Can link many private networks
Does not allow you to disable the DHCP Allocator or the DNS Proxy	Allows you to disable the DHCP Allocator or the DNS Proxy, so ICS cannot be used on a network already using a DHCP Server or DHCP Relay Agent

INSTALLING AND CONFIGURING INTERNET CONNECTION SHARING

Installing and configuring ICS is actually one of the simplest things you do in Windows. As you learned previously, though, this ease comes at the price of a good deal of flexibility. ICS is primarily for users with a small home or office network on a single network segment and a single Internet connection to share. In addition, unless you run Windows 2000 Server on the computer with the Internet connection, ICS is your only choice.

> **Note**
> This chapter focuses on using ICS in Windows 2000. Though ICS is available in Windows 98 Second Edition and Windows Millennium Edition, the configuration differs a good deal from the configuration in Windows 2000 and gives you even less control than Windows 2000. Also, the certification exam includes only the Windows 2000 version.

Installing the ICS service

You must meet only a couple of requirements before enabling ICS. First, you must make sure that the computer on which you plan to enable it (called the ICS computer from now on) actually has a functioning Internet connection, whether that connection is a 56 KB modem, cable modem, or some other type. Second, you must make sure that you have a network adapter installed in the ICS computer, that the adapter is configured and functioning properly, and that it connects properly to the other computers on the network.

When you meet these requirements, you are ready to install ICS. Hands-on Project 9-1 at the end of the chapter outlines the steps for installing ICS, but all you really need to do is open the properties dialog box for the Internet connection. (You can find it in the Network and Dial-up Connections container in the Control Panel.) Click the Sharing page, shown in Figure 9-1, and select the Enable Internet Connection Sharing for this connection option. If you want the connection to start automatically whenever other computers need to connect to the Internet (and you probably do), also select the Enable on-demand dialing option.
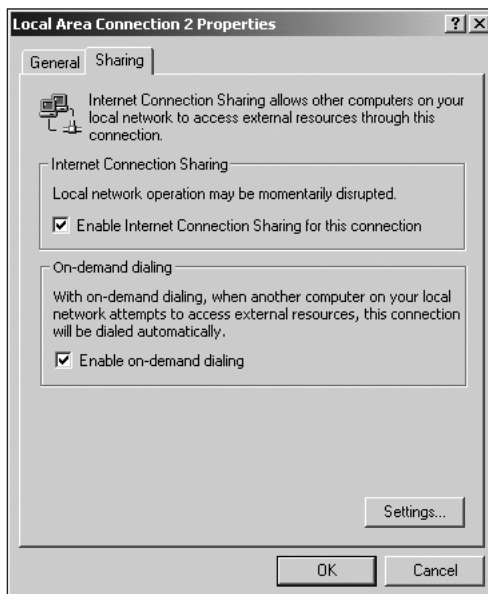
9

Figure 9-1   Sharing a connection with ICS

When you install ICS, several changes take place. These include:

- The network adapter in the ICS computer is assigned the IP address 192.168.0.1 and the subnet mask 255.255.255.0. If you recall from earlier in the chapter, this is the first address in one of the private addressing ranges.

■ The ICS service starts and is configured to start automatically each time the computer starts. You can change this behavior, as well as stop and start the service manually, using the Services Control Panel.

■ The DHCP Allocator service starts and is configured to start automatically with Windows. The allocator dynamically assigns IP addresses to other clients on the network using the IP address range 192.168.0.2 through 192.168.0.254 and the subnet mask 255.255.255.0.

Once the ICS computer is configured, you only need to ensure that all other computers on the network are configured to obtain IP addresses automatically and everything should work just fine.

Configuring ICS

With ICS enabled, configuring also takes place from the Sharing page of the adapter's properties dialog box, shown in Figure 9-1. The Settings button becomes available, and clicking it opens a dialog box that lets you configure two groups of settings that determine what entries are preloaded in the NAT mappings table on the ICS computer. Two property pages, Applications and Services, represent these groups of settings, which the next two sections discuss.

Applications Properties

The Applications page, shown in Figure 9-2, controls static outbound mappings. You use these mappings to create predefined routings for Internet services that you want users to be able to access. Normally, you do not need to worry about configuring these routings but might need to if a user's application must use a specific port number or make additional associated connections.

To add a mapping, just click the Add button to open the Internet Connection Sharing Application dialog box shown in Figure 9-3. In this dialog box, fill in the Name of application (name it anything you like), the Remote server port number and type (TCP or UDP), and the Incoming response ports that servers use to send information back to the client.
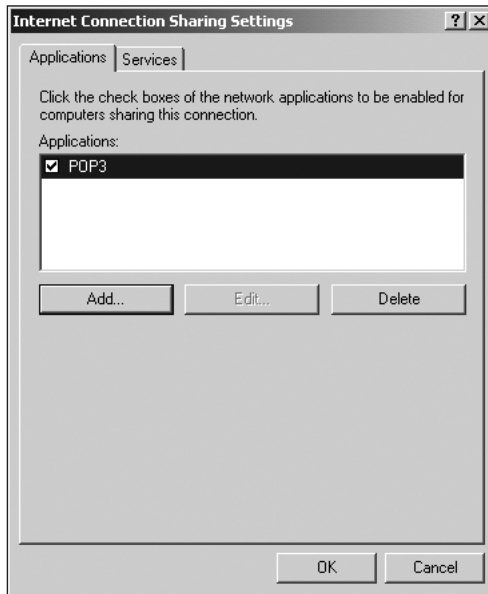
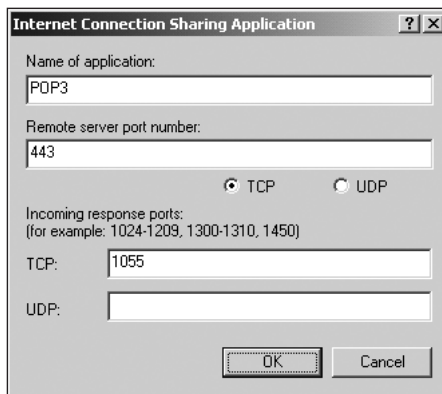**Figure 9-2**   Applications property page for ICS setting

Figure 9-3   Adding an application mapping in ICS

Services Properties

The Services page, shown in Figure 9-4, lets you control static inbound mappings. You use this feature to allow hosts on the Internet to access certain resources on the private network. Six of the most common service types are listed (but not enabled) on the page: FTP, IMAP3, IMAP4, SMTP, POP3, and TELNET.

Figure 9-4   Services property page for ICS settings

To enable a service for inbound connections, first turn it on by checking the box next to the service. Then, click the Edit button to open the Internet Connection Sharing Service dialog box shown in Figure 9-5. Note that most options are dimmed, including the Name of service, the Service port number, and the type of port. This is because these services must use the ports commonly associated with the protocols, so that outside applications can access the service without special configuration. The one setting you need to change is the name or the address of the server on the private network that hosts the service. For example, you might have a specific server dedicated to handling POP3 mail.

Figure 9-5   Enabling a service for an inbound connection

Adding a new service (using the Add button shown in Figure 9-4) uses the same dialog as editing a predefined service (Figure 9-5), but you need to enter a name and port settings in addition to a server name.

INSTALLING AND CONFIGURING NETWORK ADDRESS TRANSLATION

The NAT protocol offers much more potential for configuration than you just saw in its ICS implementation. If you run Windows 2000 Server or Advanced Server, you can implement NAT in its full glory by installing it as a routing protocol in the Routing and Remote Access snap-in. This, of course, requires that the Routing and Remote Access Server service is enabled on the server.

As with ICS, you must meet some preliminary requirements before installing NAT. First, you need to make sure that your Internet connection (or connections, since NAT supports multiple public interfaces) works. Next, you need to make sure that any adapters connected to internal networks are configured properly.

Installing the NAT Service

Once you take care of the preliminary requirements, it's time to install NAT. If you have not already configured RRAS for remote access or routing (Chapters 6 and 7 focus on these procedures), a simple wizard can guide you through the process of setting up RRAS with NAT enabled and configured for Internet sharing. Alternately, you can disable RRAS and re-enable it to remove all current settings and launch the wizard again.

If you already set up and configured RRAS and now want to add support for NAT, you do so by first ensuring that your server supports routing and then installing NAT as a routing protocol in the RRAS snap-in. Once you do this, you then add the NAT protocol to the interfaces you want to use and configure the protocol and interfaces for use. This section discusses both of these procedures.

Installing NAT Along with RRAS

If you recall from Chapters 6 and 7, RRAS is actually installed by default along with Windows 2000 Server but left disabled. You just have to enable it. This section provides an overview of the set-up process and the choices you make.

First, you must log on to the server with Administrator privileges and open the Routing and Remote Access utility from the Administrative Tools program group on the Start menu. Figure 9-6 shows this utility, which is actually a snap-in for the Microsoft Management Console used to control most management features of Windows 2000.
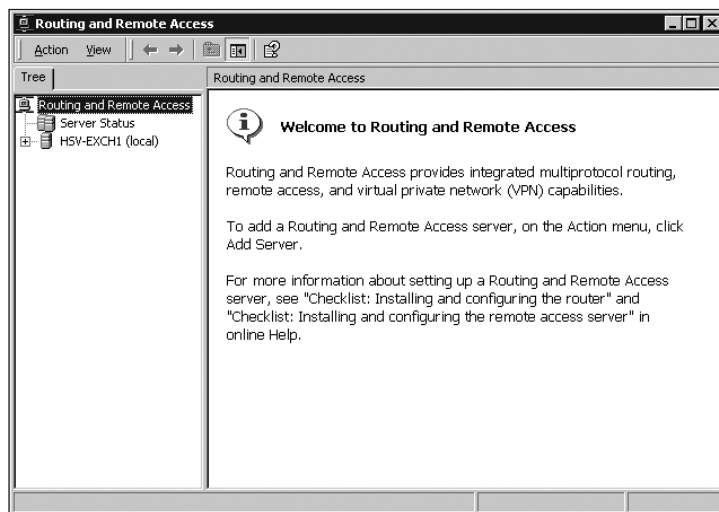
Figure 9-6 RRAS snap-in

In the tree in the left pane, find and right-click the name of the server. From the shortcut menu that appears, choose the Configure and Enable Routing and Remote Access command to begin the Routing and Remote Access Server Setup Wizard. The setup wizard takes you through several configuration steps. The first asks you to select the type of configuration you want to install. Figure 9-7 shows this screen. Choose the Internet connection server option. For details on some of the other options, see Chapter 6.


Figure 9-7 Installing RRAS as an Internet connection server

Next, the wizard asks whether you want to set up ICS or NAT. If you select ICS, a dialog box opens, telling you to use the Network and Dial–Up Connections folder to configure ICS. You do this following the procedures outlined earlier in this chapter. To set up a NAT server, of course, you must choose the NAT option.

In the next step, the wizard asks you to choose the Internet connection that you want to share, as shown in Figure 9-8. You can choose a connection from the list (you can always set up additional connections later), or you can create a new demand–dial connection. If you choose an existing connection, just pick one from the list and click Next. If you choose to create a demand-dial connection, the Demand Dial Interface Wizard opens and allows you to configure the interface before proceeding. Chapter 7 details how to set up a demand–dial interface with this wizard.
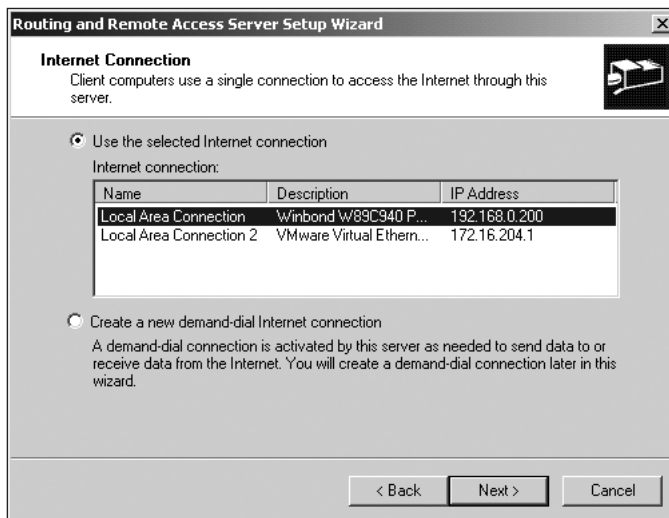

Figure 9-8   Choosing a connection to share in the RRAS Setup Wizard

Once you finish this screen, the wizard closes and the NAT server is set up. You are ready to configure the protocol or set up any additional interfaces using the RRAS snap-in. This chapter covers these procedures a bit later.

Installing NAT on an Existing RRAS Server

If you already enabled RRAS to provide remote access or routing functions, installing the NAT protocol is simple. Hands–on Project 9-2 at the end of the chapter outlines the actual steps involved. Once you install the protocol, you are ready to set up interfaces and config–ure other NAT properties.

Configuring NAT Interfaces

In earlier chapters you learned that, when working with RRAS, you must actually install and configure an interface in the RRAS snap-in before RRAS can utilize the actual network interface that the RRAS interface represents. NAT is no different. Before you can use NAT on your network, you must make sure that a **NAT interface** exists both for any interfaces on your local network and any interfaces on the public network. Following one simple rule when setting up your interfaces is best: create the interfaces for the local network first and the public network second.

Adding a NAT Interface

Adding an interface is a straightforward procedure that simply involves right-clicking the Network Address Translation container in RRAS, choosing a New Interface command, and then selecting the appropriate network adapter for which to create the interface. Hands-on Project 9-3 at the end of the chapter outlines the actual steps involved in creating a public interface. Creating a private interface follows the same procedure. Right after you create the interface, a set of property pages for the interface opens so that you can provide further configuration information. You can also open these pages later by right-clicking the interface object (shown in Figure 9-9) and choosing Properties from the shortcut menu.
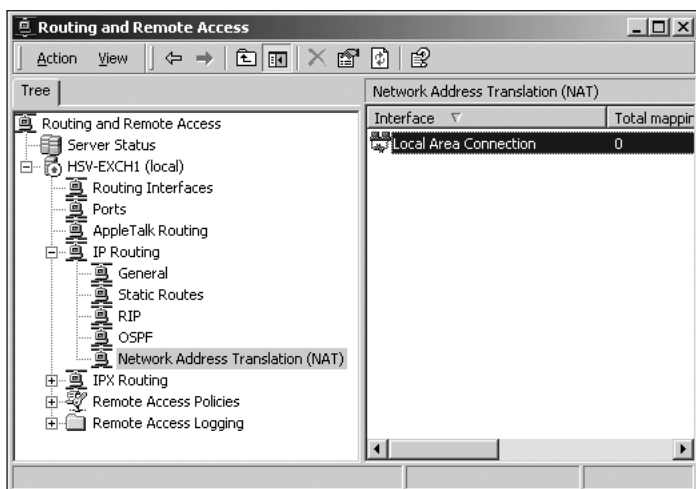

Figure 9-9   NAT Interface object in RRAS

Configuring NAT Interface Properties

Each NAT interface has its own set of property pages that is individually configurable. The three property pages for a public NAT interface are General, Address Pool, and Special Ports. The next few sections discuss each of these. The only available page for a private NAT interface is General, which is identical to the General page for the public interface.

General Properties The General page, shown in Figure 9-10, lets you choose the type of interface. You have two choices. The first is to create an interface connected to the private network. The second choice is to create an interface connected to the public network. The Translate TCP/UDP headers option controls whether the built-in NAT editors (discussed earlier in the chapter) are functional. You should always turn this on if you want computers on the private network to communicate with the outside world.
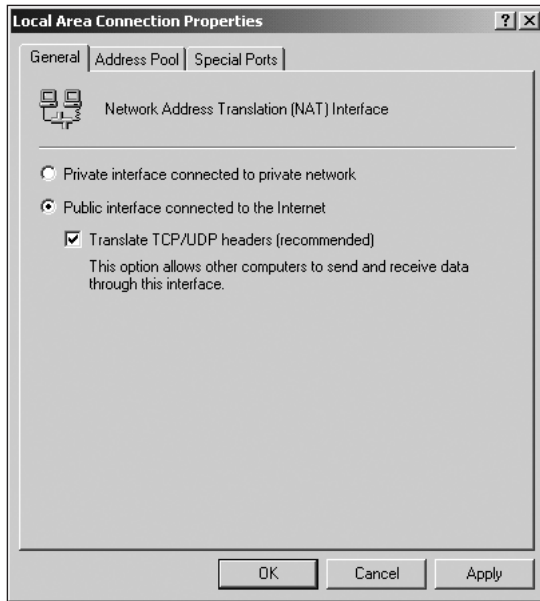

Figure 9-10   General property page of a NAT public interface

Address Pool Properties You use the Address Pool page, shown in Figure 9-11, to control the public IP addresses associated with the interface. The window lists any ranges of addresses you specified. To create a new range, click the Add button and supply the starting and ending IP addresses and the subnet mask for the range. To specify a single IP address, just enter it as the starting address and leave out the ending address.

The Reservations button lets you reserve individual IP addresses from the public range and add static mappings in the NAT table that point to particular hosts on your private network. In other words, this gives you a way to let a specific computer on your private network have a static IP address exposed to the public interface. This allows you, for example, to create a Web server and register a domain name for that Web server using the public IP address.

Special Ports Properties The Special Ports page, shown in Figure 9-12, provides another way to edit the NAT mapping table; it allows you to specify to which ports inbound traffic should map. For example, you could set it up to route all incoming traffic on port 110 (the POP3 common port) to a specific port number on a specific host on the private network—a POP3 server, most likely.
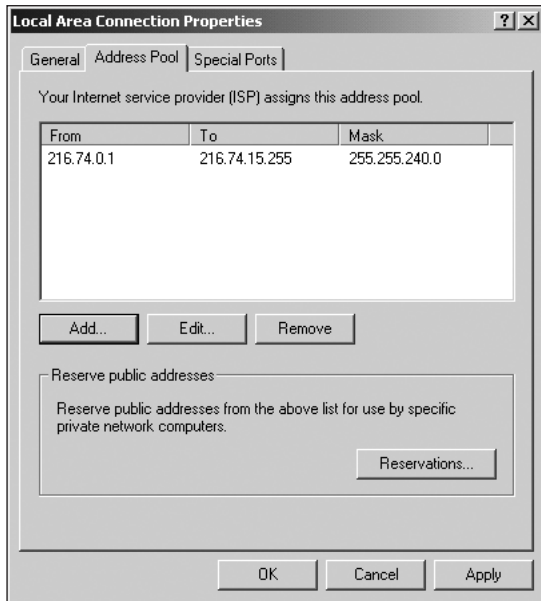
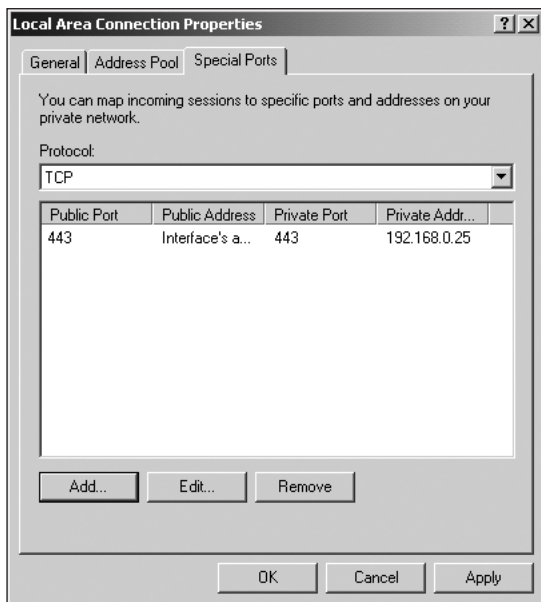**Figure 9-11**   Address Pool property page of a NAT public interface

Figure 9-12   Special Ports property page of a NAT public interface

For each protocol listed in the Protocol drop-down menu, you can specify any number of public port numbers that you want channeled to special private hosts. Just select the protocol and then use the Add button to open the Edit Special Port dialog box shown in Figure 9-13.
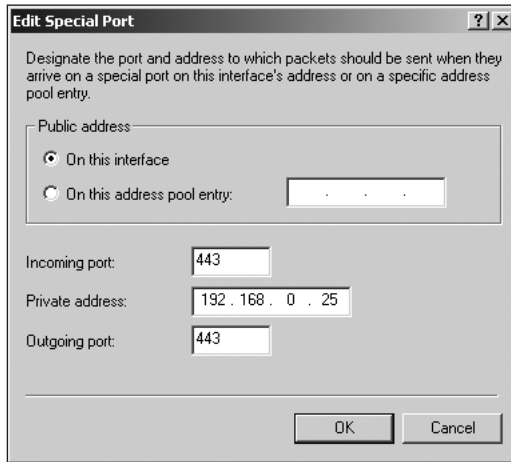

Figure 9-13   Editing a special port

This dialog box sports four controls:

- *Public Address*: controls what public address can receive traffic for the port. Choose the On this interface option (the default choice) to accept traffic on the specified port for all public IP addresses in the address pool. Choose the On this address pool entry option to specify only a specific IP address.

- *Incoming port*: specifies the port number that public hosts use to contact the service.

- *Private address*: specifies the server to which the incoming traffic should be routed.

- *Outgoing port*: specifies the port used for outbound traffic generated by hosts on the private network.

Configuring NAT Properties

In addition to setting up and configuring the individual NAT interfaces, you can set a number of global parameters for the NAT protocol itself. You can access these parameters by right-clicking the Network Address Translation container in the RRAS snap-in, shown in Figure 9-9, and choosing Properties from the shortcut menu. The four property pages for the NAT protocol are General, Translation, Address Assignment, and Name Resolution. The following sections cover each of these.

General Properties

You use the General page, shown in Figure 9-14, only to configure the level of event logging that the NAT protocol sends to the Windows 2000 system event log. The default is to log only errors, but higher levels of logging may be useful in troubleshooting problems with the protocol. You can learn more about the specific levels of logging in Chapter 6.
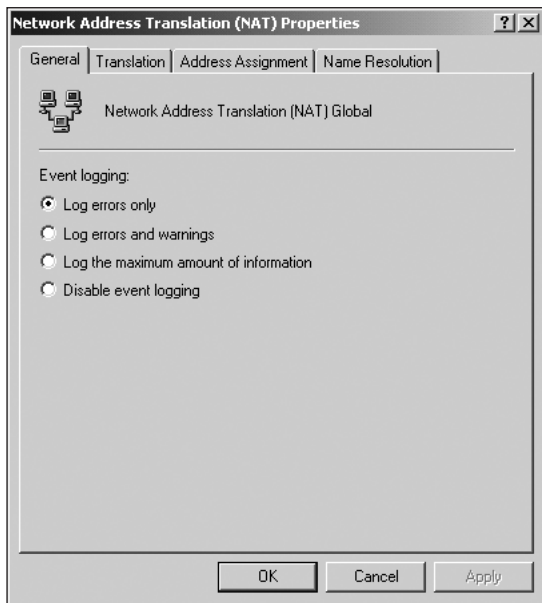

Figure 9-14   General property page of NAT Properties

Translation Properties

The Translation page, shown in Figure 9-15, lets you set the lifetime for both TCP and UDP mappings in the NAT table. The defaults are to keep TCP entries for 24 hours and to keep UDP entries for one minute; for most applications, these defaults work just fine. The Applications button opens a separate dialog box that lets you add, remove, and edit application mappings. This dialog works the same as the Applications page described for editing ICS properties earlier in the chapter and illustrated in Figure 9-2.

Address Assignment Properties

The Address Assignment page, shown in Figure 9-16, controls whether the DHCP Allocator is used or not. With this option enabled, you can specify the range of addresses the allocator can assign by entering a starting IP address and a subnet mask. By default, the same range used by ICS is used: 192.168.0.1 through 192.168.0.254. Use the Exclude button to specify IP addresses within the range that the allocator should not assign.
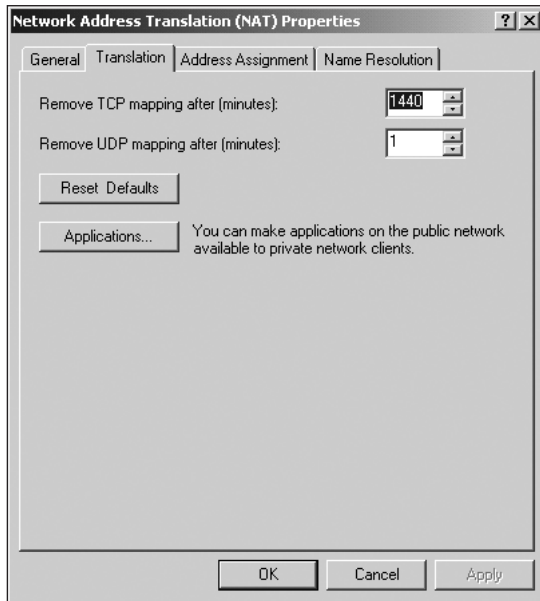
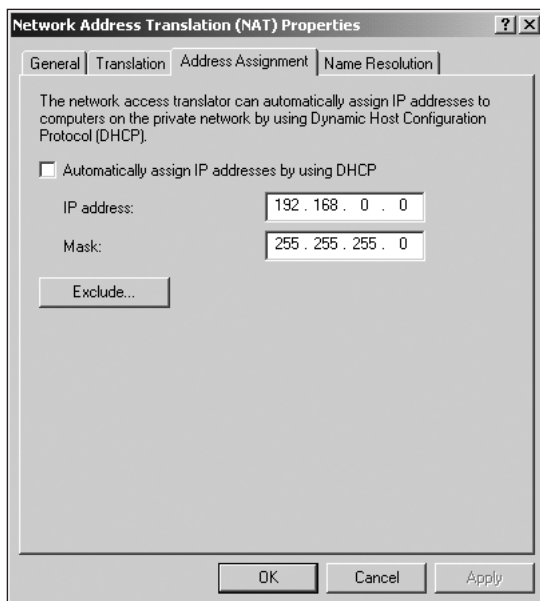**Figure 9-15** Translation property page of NAT Properties

Figure 9-16 Address Assignment property page of NAT Properties

> **Tip** If you have no other form of DHCP service on your network and you do not check the option on the Address Assignment page for the NAT protocol, NAT does not work. This is something to look out for on the certification exam and in the real world.

Name Resolution Properties

The Name Resolution page, shown in Figure 9-17, controls whether the NAT server should resolve DNS names to IP addresses for connecting clients. Enabling the Clients using Domain Name System (DNS) option activates the name resolution component of NAT and specifies the NAT server as the default DNS server for clients on the private network via the DHCP Allocator. With the option disabled, another DNS solution must be present on the network. The other option on this page, Connect to the public network when a name needs to be resolved, specifies whether a demand-dial interface is invoked just to resolve a DNS name.

Figure 9-17   Name Resolution property page of NAT Properties

CHAPTER SUMMARY

❑ The Network Address Translation protocol provides a way for multiple computers on a network to share a single connection to the Internet via an Internet Service Provider. In Windows 2000, this protocol comes in two flavors: ICS, a simplified version of the Network Address Translation protocol that is easy to configure and manage; and NAT, a full version of the protocol that is more flexible but also more difficult to set up and only available on Windows 2000 Server.

❐ A NAT server is basically an IP router that maps the IP addresses and TCP/UDP port numbers of packets as those packets are forwarded between the public and private inter-faces of the NAT server. Two forms of mapping are available in NAT. Dynamic map-pings are created when users on the private network initiate traffic with a public Internet location. Static mappings define in advance the mapping of certain addresses and ports instead of letting it happen automatically. Static mappings are required for hosting any services on the private network that will be available to the Internet.

❐ You install ICS using a single check box on the Sharing property page of an Internet connection's properties. You can configure whether demand-dialing should be used and specify some limited application and port mapping for ICS, but that's about it for config-uration. NAT is installed (on Windows 2000 Server only) as a routing protocol within the RRAS snap-in. After you install the protocol, you must create and configure any public and private interfaces you want the NAT protocol to use. You can also configure a number of properties for the protocol itself. Aside from being a good bit more config-urable than ICS, NAT offers other advantages over ICS as well. These include the ability to control the DHCP Allocator and DNS Proxy (they are always on in ICS) and the fact that NAT can maintain multiple public IP addresses while ICS can only maintain one.

9

KEY TERMS

DHCP Allocator — Simplified version of a DHCP server used by NAT to assign IP addressing information automatically to clients on the private network.

DNS proxying — Method of relaying DNS name resolution requests from clients on a private network through the NAT server to a DNS server on the Internet.

dynamic mappings — Created when users on the private network initiate traffic with a public Internet location. The NAT service automatically translates the IP address and source ports and adds these mappings to its mapping table.

Internet Connection Sharing (ICS) — Simplified version of the NAT protocol that is easy to configure and manage and is available in Windows 98, Windows Millennium Edition, Windows 2000 Server, and Windows 2000 Professional. ICS is not as config-urable as NAT.

NAT editor — Installable component that modifies packets so NAT can translate them. Windows 2000 includes built-in NAT editors for protocols, including FTP, ICMP, PPTP, and NetBT.

NAT interface — Virtual interface in the RRAS snap-in that represents an actual private or public network interface on the NAT server.

Network Address Translation (NAT) — Protocol that provides a way for multiple computers on a network to share a single connection to the Internet via an Internet Service Provider. NAT also refers to the full implementation of the protocol within the Routing and Remote Access Service in Windows 2000 Server.

private address — Any address belonging to one of the three ranges of IP addresses designated as private by Internet authorities. A host with a private address may only communicate with hosts on the Internet through a service such as NAT.

public address — Any address not belonging to one of the three ranges of IP addresses designated as private by Internet authorities.

SOHO — Acronym that stands for Small Office/Home Office. SOHO networks are considered the main beneficiaries of ICS and NAT. Though they vary a great deal in configuration, a SOHO network, as defined by Microsoft, has one network segment, uses peer-to-peer networking, and supports TCP/IP.

static mappings — Define in advance how to map certain addresses and ports instead of letting mapping happen automatically. Although you can create static mappings for outbound traffic, the most common reason to use a static mapping is to host some form of Internet service (that is, Web server, FTP server, and so forth.) on a private computer.

REVIEW QUESTIONS

1. On which of the following operating systems can the NAT protocol run?

   a. Windows 98 Second Edition

   b. Windows Millennium Edition

   c. Windows 2000 Professional

   d. Windows 2000 Server

2. Which of the following does *not* happen when you install ICS?

   a. local network adapter's IP address is reconfigured.

   b. DHCP Allocator is enabled.

   c. Internet connection is configured automatically.

   d. ICS service is configured to start automatically when Windows starts.

3. NAT must maintain mapping tables that link which of the following?

   a. Source port and address with the destination port and address

   b. Source port and address with the destination port and address of the NAT server

   c. Source port and NAT server address with the destination port and address

   d. Source port and address with the destination address and NAT server port

4. The _____ is the NAT component responsible for assigning IP addresses to local clients on the private network.

5. The ICS service assigns IP addresses ranging from 192.168.0.1 through 192.168.0.254 by default, but you can change this range if you want. True or false?

6. Which of the following protocols does not work over a NAT connection?

   a. TCP/IP

   b. IPSec

   c. FTP

   d. PPTP

7. You cannot disable the DNS proxy in NAT. True or false?

8. Which of the following must you specify when defining a NAT special port? Choose all that apply.

   a. Public address to receive traffic for the port

   b. Port numbers used for inbound and outbound traffic

   c. Private IP address that receives traffic on the special port

   d. Subnet mask used for the port

9. If you are using a modem rather than a dedicated link to the Internet, which option must you enable?

   a. ICS automatic dialing

   b. On-demand dialing

   c. Automatic dialing

   d. Dynamic linking

10. RRAS can act as a NAT server and a remote access server simultaneously. True or false?

11. A _____ is used to support the translation of traffic generated by protocols or applications that do not use TCP or UDP.

12. For what is the Translation property page of the NAT protocol used?

    a. To create application-specific port mappings

    b. To specify which NAT editor to use

    c. To create port mappings for individual hosts

    d. To specify which port filters to apply

13. To allow a host on your private network to act as a Web server accessible from the Internet, you must configure a _____.

14. You must decide whether to use NAT or ICS on your small office network. You want to choose the simplest service to set up and manage, but you do need to run an FTP server inside your private network and make it accessible to users on the Internet. Which service would you choose?

15. A _____ is an automatic translation of IP addresses and source ports performed by the NAT protocol when users on the private network initiate traffic with a public Internet location.

16. You have a small network with two network segments and want to keep different subnet addresses for them. How do you do this?

    a. Add NAT interfaces for both networks.

    b. Disable the DHCP Allocator.

    c. Define two static address pools with the subnets you want to use.

    d. Manually assign IP addresses to the server's internal interfaces.

9

17. You use the _____ property page of the NAT interface to choose whether it is a public or private interface.

18. You can use ICS only with demand–dial connections. To use a dedicated Internet connection, you must configure NAT. True or false?

19. Which of the following IP addresses are private addresses?

    a. 10.35.202.1

    b. 172.16.18.2

    c. 172.101.201.44

    d. 192.168.201.1

20. Remote Procedure Calls can be used over NAT. True or false?

HANDS-ON PROJECTS

All Hands-on Projects in this chapter require at least one server computer set up as described in the lab set-up section in the front of this book.

Project 9-1

To install Internet Connection Sharing, you must log on to the local computer under an account with Administrator privileges.

To install ICS on a local computer:

1. Click **Start**, point to **Settings**, and then click **Network and Dial–up Connections**.

2. Right-click the icon for the adapter that represents your Internet connection, and select the **Properties** command.

3. Click the **Sharing** tab to switch to that page.

4. Select the **Enable Internet Connection Sharing For This Computer** option.

5. Select the **Enable On-Demand Dialing** option.

6. Click the **OK** button.

7. A dialog box appears, warning you that the IP address of the adapter will change if you continue. Click **Yes** to finish the installation.

Project 9-2

To install NAT on an existing RRAS Server:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then select **Routing and Remote Access**.

2. Find the server you want to configure in the left pane, and expand it.

3. Inside the **IP Routing** container for the server, right-click the **General** container and select the **New Routing Protocol** command from the shortcut menu.

4. In the **New Routing Protocol** dialog box that opens, select the **Network Address Translation** item from the list of routing protocols and click **OK**. The **IP Routing** container should now contain a new object named **Network Address Translation**.

Project 9-3

To add and configure a public NAT interface:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then select **Routing and Remote Access**.

2. Find the server you want to configure in the left pane, and expand it.

3. Inside the **IP Routing** container for the server, right-click the **Network Address Translation** container and select the **New Interface** command from the shortcut menu.

4. In the **New Interface for Network Address Translation** dialog box, select the adapter you want to use for the interface from the list and click **OK**.

5. The **Network Address Translation Properties** dialog box appears. On the **General** page, select the **Public interface connected to the Internet** option and click **OK**.

9

CASE PROJECTS

Case 1

Your small network consists of two subnets. You configured one subnet with the network ID 192.168.0.0 and the other with the network ID 192.168.1.0. A computer running Windows 2000 Server and configured with RRAS serves as a router between the two networks. All computers on both network segments are configured with static IP addressing. You just installed a DSL line and successfully established an Internet connection from the Windows 2000 Server. Describe the steps you must take in order to share that Internet connection with both network segments.

Case 2

You provide consulting services for a small company with a single network segment and 12 computers, all running Windows 2000 Professional. You just helped the company install a cable modem, and the owner wants all computers on the network to have access to the connection. The owner has read about NAT and is convinced that he needs to install a Windows 2000 Server computer and configure it with RRAS. His main reason: he and a few employees want to connect to the network from home using Virtual Private Networking. Write an explanation detailing why ICS would meet his needs and why it would be preferable over NAT.